arXiv:1604.03499v1 [math.CA] 12 Apr 2016

NOISY 1-BIT COMPRESSED SENSING EMBEDDINGS ENJOY A RESTRICTED ISOMETRY PROPERTY

SCOTT SPENCER

Abstract. We investigate the sign-linear embeddings of 1-bit compressed sensing given by Gaussian measurements. One can give short arguments concerning a Restricted Isometry Property of such maps using *Vapnik-Chervonenkis dimension* of sparse hemispheres. This approach has a natural extension to the presence of additive white noise prior to quantization. Noisy one-bit mappings are shown to satisfy an RIP when the metric on the sphere is given by the noise.

1. Introduction

Compressed sensing is a modern data processing scheme that is proving useful in many scientific areas, such as MR imaging, radar, astronomy: see [1, 2, 10] for more details. The overarching goal is to reconstruct a *signal* $x \in \mathbb{R}^n$ from the *measurements* $Ax \in \mathbb{R}^m$ ($m \ll n$) given the sensing matrix $A \in \mathbb{R}^{m\times n}$ and some constraint on the set of signals. Without such a constraint, this is an ill-posed inverse problem, while more information about the signal x may make the objective approachable. One common situation is that the signal is sparse: for a signal $x = (x^1, \dots, x^n)$, we say x is s-sparse if $|\{x^j \neq 0\}| \leq s$. A successful program for reconstructing sparse signals is ℓ_1-minimization. This convex optimization algorithm is tractable and perfectly reconstructs s-sparse vectors (and well approximates them in the presence of noise) if the sensing matrix A has the (s, δ)-RIP with small enough δ [5]. A matrix A is said to have the (s, δ)-RIP if

$$(1-\delta)\|x-y\|_2^2 \leq \|Ax - Ay\|_2^2 \leq (1+\delta)\|x-y\|_2^2$$

for all pairs x, y of s-sparse vectors. The object of our interest is the analogue, i.e., dimension reducing quasi-isometric embeddings of sparse vectors, in the 1-bit sensing framework.

1.1. 1-Bit Sensing.

We study the dimension reducing *sign-linear* maps of 1-bit compressed sensing. Associated to each $A \in \mathbb{R}^{m\times n}$ is the sign-linear map

$$\Phi_A : \mathbb{S}^{n-1} \to \mathcal{H}^m \tag{1.1}$$
$$\Phi_A x = \operatorname{sgn}(Ax),$$

where $\mathcal{H}^m$ is the Hamming Cube $\{\pm 1\}^m$, the *sgn* map is applied component-wise, and

$$\operatorname{sgn}(x) = \begin{cases} +1, & x > 0 \\ -1, & x \leq 0. \end{cases}$$

We restrict our attention to the sphere since any two signals that differ only in norm will have identical measurements. In the larger realm of compressed sensing, 1-bit sensing is the

2010 *Mathematics Subject Classification.* 94A12; 60G15; 94A24.

Key words and phrases. quantization, dimension reduction, RIP, quasi-isometry, binary stable embedding, VC-dimension.

Research supported in part by grant NSF-DMS 1265570.

case of extreme quantization: only the sign-bit of each linear measurement is preserved. The concept was initially suggested by Boufounos-Baraniuk [4] in 2008.

Let $\mathbb{S}_s^{n-1}$ denote the set of n-dimensional, unit length s-sparse signals. The (s,δ)-*Restricted Isometry Property*, or (s,δ)-RIP, analogue for Φ_A that we investigate is

$$\sup_{x,y\in\mathbb{S}_s^{n-1}} |d_{\mathcal{H}^m}(\Phi_A x, \Phi_A y) - d(x,y)| \leq \delta,$$

where $d(\cdot,\cdot)$ is geodesic distance on the sphere, and $d_{\mathcal{H}^m}(\cdot,\cdot)$ is the Hamming metric:

$$d_{\mathcal{H}^m}(a,b) := \tfrac{1}{m}\left|\{1 \leq k \leq m : a_k \neq b_k\}\right|.$$

The reader may notice that the 1-bit RIP given above is single-scale, while the original RIP is multiscale. This modification is unavoidable; given $A \in \mathbb{R}^{m\times n}$ and $\epsilon > 0$, there are $x, y \in \mathbb{S}_s^{n-1}$ such that $d(x,y) \leq \epsilon$ and $d_{\mathcal{H}^m}(\Phi_A x, \Phi_A y) \geq \frac{1}{m}$. This formulation of the RIP has been studied theoretically, see [3, 12]; it also plays a role in sparse signal recovery from 1-bit measurements, e.g. [8, 9].

It is natural to consider the effects of noise on a 1-Bit embedding. We consider the case of additive white noise prior to quantization. When we consider a random sensing matrix A and random noise η, we always assume they are independent. Associated to a matrix $A \in \mathbb{R}^{m\times n}$ and vector $\eta \in \mathbb{R}^m$ is a 1-bit embedding of the form

$$\begin{aligned} \Phi_A^\eta &: \mathbb{R}^n \to \mathcal{H}^m \\ \Phi_A^\eta x &= \operatorname{sgn}(Ax + \eta). \end{aligned} \tag{1.2}$$

Taking $\eta \sim \mathcal{N}(\mathbf{0}, \sigma^2 I_m)$, the (s,δ)-RIP analogue for Φ_A^η that we investigate is

$$\sup_{x,y\in\mathbb{S}_s^{n-1}} |d_{\mathcal{H}^m}(\Phi_A^\eta x, \Phi_A^\eta y) - d^\sigma(x,y)| \leq \delta.$$

The distance $d^\sigma(\cdot,\cdot)$ is a distorted version of the geodesic distance, defined in (2.1) and discussed in Section 5. The affects of the additive white noise on the RIP are analyzed by increasing the Gaussian measurements' dimension by one and lifting the sphere to one higher dimension by padding with σ^2.

2. Outline and Main Results

For $x \in \mathbb{S}^{n-1}$, set $H_x = \{p \in \mathbb{S}^{n-1} : \langle p, x\rangle > 0\}$, the hemisphere associated to x. Denote by $H^{n,s}$ the set of hemispheres of $\mathbb{S}^{n-1}$ associated to s-sparse signals: $H^{n,s} = \{H_x : x \in \mathbb{S}_s^{n-1}\}$. The first result listed here gives a useful upper bound on the VC-dimension, defined in Section 3.1, of $H^{n,s}$. The result easily applies to half-spaces, a well studied classification scheme in learning theory; it is well known that the VC-dimension of half-spaces in $\mathbb{R}^n$ indexed by s-sparse vectors is $\mathcal{O}(s\log n)$. The theorem below is slightly better, but we are unsure if it is known. We include the proof in Section 3.2 for completeness, and note that it is quite surprising to find the popular $s\log(n/s)$ quantity. Throughout, $x \lesssim y$ means there is an absolute $C > 0$ such that $x \leq Cy$.

Theorem 2.1. *$VC(H^{n,s}) \lesssim s\log(n/s)$.*

Definition 2.2. *Let $\Phi : \mathbb{S}_s^{n-1} \to \mathcal{H}^m$. We say Φ has the (s,δ)-RIP if*

$$\sup_{x,y\in\mathbb{S}_s^{n-1}} |d_{\mathcal{H}^m}(\Phi x, \Phi y) - d(x,y)| \leq \delta.$$

Of note in Definition 2.2 is the metric $d(\cdot,\cdot)$, which is **not** the euclidean distance, but rather the geodesic distance on the sphere, normalized so that antipodal points are unit distance apart:

$$d(x,y) := \tfrac{1}{\pi}\arccos\big(\langle x,y\rangle\big).$$

This choice of metric is natural since it is the expectation of $d_{\mathcal{H}^m}(\Phi_A x, \Phi_A y)$.

In Section 4 we employ a standard entropy integral argument to bound a supremum, indexed by pairs of s-sparse vectors. This is an alternative proof of a recent result of Bilyk-Lacey, the case of sparse vectors in [3, Theorem 1.14], which is:

Theorem 2.3. *Let $A \in \mathbb{R}^{m\times n}$ with rows drawn independently from the standard Gaussian distribution. Then for any $0<\varepsilon,\delta<1$ and $1\le s<n$, Φ_A has the (s,δ)-RIP with probability at least $1-\varepsilon$ provided*

$$m \gtrsim \delta^{-2}\left[\log(2/\varepsilon) + s\log(n/s)\right].$$

The next theorem, proved in Section 5, is the import of the paper. We consider the 1-bit sign-linear maps with additive white noise prior to quantization. A curious detail about the result is that the error due to noise is not naturally expressed in the distortion parameter, nor the number of measurements or probability of success, but rather in the metric on the sphere. That is, if the sphere is endowed with a certain "distorted" geodesic metric (2.1), the noisy embedding has the (s,δ)-RIP with the same order of measurements and probability of success as determined in Theorem 2.3. Before stating the theorem, we define:

$$d^\sigma(x,y) := \tfrac{1}{\pi}\arccos\left(\tfrac{\langle x,y\rangle+\sigma^2}{1+\sigma^2}\right) \quad \text{for } x,y\in\mathbb{S}^{n-1}. \tag{2.1}$$

We also define the following noisy version of the 1-bit RIP:

Definition 2.4. *Let $\Phi : \mathbb{S}_s^{n-1} \to \mathcal{H}^m$. We say Φ has the (s,δ,σ)-RIP if*

$$\sup_{x,y\in\mathbb{S}_s^{n-1}} |d_{\mathcal{H}^m}(\Phi x,\Phi y) - d^\sigma(x,y)| \le \delta.$$

Theorem 2.5. *Let $A\in\mathbb{R}^{m\times n}$ with rows drawn independently from the standard Gaussian distribution and $\eta\sim\mathcal{N}(\boldsymbol{0},\sigma^2 I_m)$. Then for any $0<\varepsilon,\delta<1$ and $1\le s<n$, Φ_A^η has the (s,δ,σ)-RIP with probability at least $1-\varepsilon$ provided*

$$m \gtrsim \delta^{-2}\left[\log(2/\varepsilon) + s\log(n/s)\right].$$

Remark. It is a common goal in noisy compressive sensing to "eliminate" the noise. That is, one wishes to take enough measurements so that the noise is practically negligible. Theorem 2.5 demonstrates that this possibility is controlled by the variance in the Gaussian noise model. The empirical process of interest approaches the distorted metric d^σ, which is a deterministic object that necessarily deviates from the geodesic metric when $\sigma^2>0$.

We conclude with Section 5.2, comparing the geodesic distance with the metric defined in (2.1). A crude upper bound on their difference gives a lower bound on the number of Gaussian measurements needed for a noisy embedding to have the RIP into the Hamming cube with the *geodesic* metric prescribed to the sphere. While this result is appealing for obvious reasons, Theorem 2.5 may be more useful in practice, allowing the reader to appeal to the fact that the two metrics are indeed *very* close at small scales.

Corollary 2.6. *Let $A \in \mathbb{R}^{m\times n}$ with rows drawn independently from the standard Gaussian distribution and $\eta \sim \mathcal{N}(\boldsymbol{0}, \sigma^2 I_m)$. Then for any $0 < \varepsilon < 1$ and $\delta > 1 - \frac{1}{\pi}\arccos\left(\frac{\sigma^2-1}{\sigma^2+1}\right)$, Φ_A^η has the (s,δ)-RIP with probability at least $1-\varepsilon$ provided*

$$m \gtrsim \left[\delta + \tfrac{1}{\pi}\arccos\left(\tfrac{\sigma^2-1}{\sigma^2+1}\right) - 1\right]^{-2} \left[\log(2/\varepsilon) + s\log(n/s)\right].$$

3. The VC-Dimension of Sparse Hemispheres

3.1. **VC Dimension.** Let X be a set and $\mathcal{C}$ be a collection of subsets of X. Denote by $\binom{X}{k}$ the set of subsets of X with k elements. For each $k \in \mathbb{N}$, define

$$m^{\mathcal{C}}(k) := \max_{B\in\binom{X}{k}} \left|\{B \cap C : C \in \mathcal{C}\}\right|.$$

Clearly $m^{\mathcal{C}}(k) \le 2^k$. The *Vapnik-Chervonenkis dimension* ($VC-$dimension) of $\mathcal{C}$, denoted $VC(\mathcal{C})$, is the largest integer d (if it exists) such that $m^{\mathcal{C}}(d) = 2^d$, and $VC(\mathcal{C}) = \infty$ otherwise. Alternatively, we say $\mathcal{C}$ *shatters* B if every subset of B is realized as the intersection of B with an element of $\mathcal{C}$. Then $VC(\mathcal{C})$ is the cardinality of the largest subset it shatters. For example, if $X = \mathbb{R}$ and $\mathcal{C} = \{(-\infty, t] : t \in \mathbb{R}\}$, then $VC(\mathcal{C}) = 1$; if $\mathcal{C} = \{[a,b] : a < b \in \mathbb{R}\}$, then $VC(\mathcal{C}) = 2$. VC dimension measures, in an intuitive sense, the complexity of a class of subsets.

The following lemma is a fundamental result in VC theory, and we will use it several times. A proof of the lemma and other details on the subject can be found in [6].

Lemma 3.1 (Sauer's Lemma)**.** *Let $\mathcal{C}$ be a class of subsets with $VC(\mathcal{C}) = d < \infty$. Then for any $k \ge d$,*

$$m^{\mathcal{C}}(k) \le \left(\tfrac{ek}{d}\right)^d.$$

For a class of functions $\mathcal{F} \subset \{f : X \to \{0,1\}\}$, denote by $\mathcal{C}_{\mathcal{F}}$ the set of subgraphs of functions in $\mathcal{F}$: $\mathcal{C}_{\mathcal{F}} = \{\{(x,t) : t \le f(x)\} : f \in \mathcal{F}\}$. The VC dimension of $\mathcal{F}$ is defined as $VC(\mathcal{C}_{\mathcal{F}})$, where this last quantity is the VC-dimension of a class of subsets of $X \times \mathbb{R}$. It is worth noting that if $\mathcal{F}$ is the set of indicators of subsets in the class $\mathcal{C}$, $\mathcal{F} = \{1_C : C \in \mathcal{C}\}$, then $VC(\mathcal{F}) = VC(\mathcal{C})$.

It is well known in learning theory that empirical processes in the form of (4.2) can be bounded via the VC-dimension of the indexing class. Such results are often eponymously referred to as the "VC inequality" after Vapnik and Chervonenkis, the pioneers of the theory. In Section 4 we use a version of the VC inequality from [11], which extends the VC inequality to a more general case, when a class satisfies *uniform entropy bounds*. For a function f and a probability $\mathbb{P}$, denote by $\mathbb{P}f$ the expectation $\int f\, d\mathbb{P}$. For a class of binary functions $\mathcal{F}$ and a probability $\mathbb{P}$, the packing number $D(\mathcal{F}, t, \mathbb{P})$ is the cardinality of the largest subset $\mathcal{F}' \subset \mathcal{F}$ such that $\mathbb{P}|f - g| > t^2$ for all $f \ne g \in \mathcal{F}'$. Finally, set

$$D(\mathcal{F}, t) := \sup_{\mathbb{P}} D(\mathcal{F}, t, \mathbb{P}),$$

where the supremum is taken over all discrete probabilities. Then [11, corollary 1] reads

Theorem 3.2. *Suppose*

$$\int_0^\infty \sqrt{\log D(\mathcal{F}, t)}\, dt < \infty.$$

Then there exists an absolute constant $K > 0$ such that for any $u > 0$ with probability at least $1 - 2e^{-u}$ for all $f \in \mathcal{F}$:

$$\sum_{k=1}^{m} (\mathbb{P}f - f(x_k)) \leq K\sqrt{m}\left(\sqrt{u\mathbb{P}f} + \int_0^{\sqrt{\mathbb{P}f}} \sqrt{\log(D(\mathcal{F}, t))}dt\right).$$

3.2. **Main VC Estimate.** This section is dedicated to the proof of Theorem 2.1. We begin by computing the VC-dimension of all hemispheres, the case when $s = n$.

Lemma 3.3. $VC(H_s^s) = s$.

Proof. We first observe H_s^s shattering the standard basis vectors $B = \{e_1, \ldots, e_s\}$, and hence $VC(H_s^s) \geq s$. Let $S \subset [s]$ and $B(S) = \{e_j : j \in S\}$. Define $p = (p_1, \ldots, p_s)$ by setting $p_j = 1_S(j) - 1_{S^c}(j)$. Then $B(S) = B \cap H_p$.

On the other hand, let $X = \{x_1, \ldots, x_{s+1}\}$ be an arbitrary $(s+1)$-subset of $\mathbb{S}^{s-1}$. Without loss of generality, assume

$$x_{s+1} = \sum_{k=1}^{s} \alpha_k x_k.$$

Set $A := \{x_k : \alpha_k < 0\} \cup \{x_{s+1}\}$; we'll see that for all $p \in \mathbb{R}^s$, $A \neq X \cap H_p$. For any p such that $\langle p, x_k\rangle > 0$ if $\alpha_k < 0$ and $\langle p, x_k\rangle \leq 0$ if $\alpha_k \geq 0$,

$$\begin{aligned}\langle p, x_{s+1}\rangle &= \sum_{k=1}^{s} \alpha_k \langle p, x_k\rangle \\ &= \sum_{k:\alpha_k<0} \alpha_k\langle p, x_k\rangle + \sum_{k:\alpha_k\geq 0} \alpha_k\langle p, x_k\rangle \\ &\leq 0.\end{aligned}$$

Therefore H_s^s doesn't shatter X, so $VC(H_s^s) < s + 1$. □

We are now ready to estimate $VC(H^{n,s})$. Let $d = VC(H^{n,s}) \leq n$ and choose a subset $X = \{x^1, \ldots, x^d\}$ of $\mathbb{S}^{n-1}$ shattered by $H^{n,s}$. Fix an index set $S \in \binom{[n]}{s}$; for $x \in \mathbb{S}^{n-1}$ let $x_S = \sum_{j\in S}\langle x, e_j\rangle e_j$. For any $B \subset \mathbb{S}^{n-1}$, let $B_S = \{b_S/\|b_S\| : b \in B \text{ and } b_S \neq 0\}$. Notice that $|X_S| \leq d$, so by Lemmas 3.3 and 3.1,

$$\left|\{X_S \cap H_p : p \in \mathbb{S}_S^{n-1}\}\right| \leq \left(\tfrac{ed}{s}\right)^s.$$

The natural map $\{X \cap H_p : p \in \mathbb{S}_S^{n-1}\} \to \{X_S \cap H_p : p \in \mathbb{S}_S^{n-1}\}$ via $A \mapsto A_S$ is well-defined and surjective since $\operatorname{sgn}(\langle x, p\rangle) = \operatorname{sgn}(\langle x_S, p\rangle)$ for all $x \in X$ and $p \in \mathbb{S}_S^{n-1}$. This map is also injective. Suppose $A = X \cap H_p$ and $B = X \cap h_{p'}$ are distinct, for instance $a \in A \setminus B$ (hence $a_S \neq 0$). If $a_S/\|a_S\| \in B_S$, then there is $b \in B$ such that $a_S/\|a_S\| = b_S/\|b_S\|$. But then $\operatorname{sgn}(\langle a, p'\rangle) = \operatorname{sgn}(\langle b, p'\rangle)$, a contradiction.

It follows that

$$\left|\{X \cap H_p : p \in \mathbb{S}_S^{n-1}\}\right| \leq \left(\tfrac{ed}{s}\right)^s,$$

and by the union bound, $2^d \leq \binom{n}{s}\left(\frac{ed}{s}\right)^s$. After applying a familiar version of Stirlings approximation, $\binom{n}{s} \leq \left(\frac{en}{s}\right)^s$, and some algebraic manipulation, we arrive at the inequality:

$$-\log(2)\tfrac{d}{s}e^{-\log(2)\frac{d}{s}} \leq -\tfrac{\log(2)s}{e^2 n}. \tag{3.1}$$

To simplify further, we use the lower branch of the *Lambert W function*, which is defined on $(\frac{-1}{e}, 0)$ by the relation $W_{-1}(x)e^{W_{-1}(x)} = x$. That is, W_{-1} is the inverse of the map $x \mapsto xe^x$ restricted to $(-\infty, -1)$. We use the following lower bound of W_{-1} to simplify (3.1).

Lemma 3.4. *For all* $-1/e < x < 0$, $W_{-1}(x) \geq \log(x^2)$.

Proof. Notice that W_{-1} is decreasing, as is its inverse $W_{-1}^{-1}(x) = xe^x$. Applying W_{-1}^{-1} to each side of the equation in the statement and dividing by x, we find the equivalent: $x\log(x^2) \leq 1$ for all $\frac{-1}{e} < x < 0$. This holds since $x \mapsto x\log(x^2)$ is decreasing on $(-1/e, 0)$ and $(\frac{-1}{e})\log(\frac{1}{e^2}) = \frac{2}{e} < 1$. □

Applying the decreasing W_{-1} to both sides of (3.1) and using Lemma 3.4 gives:

$$d \leq \tfrac{2}{\log(2)} s \log\left(\tfrac{ne^2}{s\log(2)}\right).$$

4. The RIP of 1-Bit Embeddings

This section proves Theorem 2.3. Let $A \in \mathbb{R}^{m\times n}$ with rows $\{g_k\}_{k=1}^m$ drawn independently from the standard Gaussian distribution $\mathcal{N}(\mathbf{0}, I_n)$. The Hamming distance between the images of two signals x and y under the 1-Bit embedding Φ_A is

$$d_{\mathcal{H}^m}(\Phi_A x, \Phi_A y) = \frac{1}{m}\sum_{k=1}^{m} \frac{1 - \operatorname{sgn}\langle x, g_k\rangle \operatorname{sgn}\langle y, g_k\rangle}{2}.$$

For $x, y \in \mathbb{S}^{n-1}$ we call $W_{x,y} := H_x \triangle H_y$ (the symmetric difference of the two hemispheres) the *wedge* associated to x and y. Notice $\operatorname{sgn}\langle x, g_k\rangle \neq \operatorname{sgn}\langle y, g_k\rangle$ if and only if g_k is in the wedge $W_{x,y}$. The Hamming distance above can be reformulated as

$$d_{\mathcal{H}^m}(\Phi_A x, \Phi_A y) = \frac{1}{m}\sum_{k=1}^{m} 1_{W_{x,y}}(g_k).$$

The empirical processes framework suggests the sphere should be endowed with the distance $(x, y) \mapsto \mathbb{P}(W_{x,y})$. Fix $x, y \in \mathbb{S}^{n-1}$, let $g \sim \mathcal{N}(\mathbf{0}, I_n)$, and let $Z = (\langle x, g\rangle, \langle y, g\rangle)^\top$; then $Z \sim \mathcal{N}(\mathbf{0}, \Sigma)$ with

$$\Sigma = \begin{bmatrix} 1 & \langle x, y\rangle \\ \langle x, y\rangle & 1 \end{bmatrix}.$$

It is a basic computation to find

$$\begin{aligned} \mathbb{P}(W_{x,y}) &= \tfrac{1}{\pi\sqrt{1-\langle x,y\rangle^2}} \int_0^\infty \int_0^\infty \operatorname{Exp}\left(\tfrac{2uv\langle x,y\rangle - u^2 - v^2}{2-2\langle x,y\rangle^2}\right) du\, dv \\ &= \tfrac{1}{\pi}\arccos\big(\langle x, y\rangle\big). \end{aligned} \tag{4.1}$$

This last quantity is the geodesic distance on the sphere that we denote by $d(x, y)$. This brings our attention to the following object:

$$\sup_{x,y\in\mathbb{S}_s^{n-1}} \left| \frac{1}{m}\sum_{k=1}^{m} 1_{W_{x,y}}(g_k) - d(x, y) \right|. \tag{4.2}$$

The above formulation is paraphrased from [3]; this is the point at which our argument deviates. To utilize the VC theory for hemispheres developed in the previous section, we bound the VC-dimension of the class of "sparse wedges" $\mathcal{W}^{n,s} := \{W_{x,y} : x, y \in \mathbb{S}_s^{n-1}\}$.

Lemma 4.1. *Let $\mathcal{C}$ be a class of subsets of X with $VC(\mathcal{C}) = d < \infty$. Let $\mathcal{C}\triangle\mathcal{C} = \{C\triangle C' : C, C' \in \mathcal{C}\}$. Then $VC(\mathcal{C}\triangle\mathcal{C}) \leq 10d$.*

Proof. Let $B \subset X$ of size $m := |B|$ to be prescribed later. For a fixed pair $C, C' \in \mathcal{C}$, Notice that

$$B \cap (C\triangle C') = [(B \cap C) \setminus (B \cap C')] \cup [(B \cap C') \setminus (B \cap C)] .$$

That is, $B \cap (C\triangle C')$ is determined by $B \cap C$ and $B \cap C'$. By Lemma 3.1, there are no more than $\left(\frac{em}{d}\right)^{2d}$ such pairs. Taking $m \geq 10d$ yields $\left(\frac{em}{d}\right)^{2d} < 2^m$. □

Along with Theorem 2.1, this lemma implies $VC(\mathcal{W}^{n,s}) \lesssim s\log(n/s)$. We use this VC-dimension estimate to bound the *packing numbers* of the sparse wedges, $D(\mathcal{W}^{n,s}, \epsilon, \mathbb{P})$, which is the largest d so that there exists $w_1, \ldots, w_d \in \mathcal{W}^{n,s}$ with $\mathbb{P}(w_i\triangle w_j) > \epsilon^2$ for all $i \neq j$. General results bounding packing numbers via VC-dimension are well-known and the argument is standard; we include a proof in the current context for completeness.

Proposition 4.2. *For $0 < \epsilon < 1$,*

$$D(\mathcal{W}^{n,s}, \epsilon, \mathbb{P}) \lesssim \left(\tfrac{1}{\epsilon^2}\right)^{VC(\mathcal{W}^{n,s})+1} .$$

Proof. Fix $0 < \epsilon < 1$. Let $d = D(\mathcal{W}^{n,s}, \epsilon, \mathbb{P})$ and let $w_1 \ldots, w_d$ such that $\mathbb{P}(w_i\triangle w_j) > \epsilon^2$ for all $i \neq j$. Let $\{X_k\}_{k=1}^n$ be independent and identically distributed on the sphere with law $\mathbb{P}$, where n will be determined later. Notice that $w_i \cap \{X_k\} \neq w_j \cap \{X_k\}$ if and only if $(w_i\triangle w_j) \cap \{X_k\}$ is nonempty. Thus the probability that there is $i \neq j$ such that $w_i \cap \{X_k\} = w_j \cap \{X_k\}$ is no more than

$$\begin{aligned}
\tbinom{d}{2} \max_{1\leq i\neq j\leq d} \mathbb{P}\left(w_i \cap \{X_k\} = w_j \cap \{X_k\}\right) &= \tbinom{d}{2} \max_{1\leq i\neq j\leq d} \left(1 - \mathbb{P}(w_i\triangle w_j)\right)^n \\
&< \tbinom{d}{2}(1-\epsilon^2)^n \\
&< d^2 e^{-n\epsilon^2} \\
&= e^{2\log(d) - n\epsilon^2}.
\end{aligned}$$

Now we take $n = \left\lceil \frac{2\log(d)+1}{\epsilon^2} \right\rceil$ so the above probability is less than one, hence there is a deterministic $X = \{x_k\}_{k=1}^n$ so that $d = |\{w_j \cap X : j \in [d]\}|$. Let $v = VC(\mathcal{W}^{n,s})$. Employing Lemma 3.1, there is $K_v > 0$ such that

$$d \leq K_v \left(\tfrac{2\log(d)+2}{\epsilon^2}\right)^v .$$

Choose d_0 large enough so that for $d > d_0$, $(2\log(d) + 2)^{v+1} < d^{1/v}$. This yields

$$d \leq \max\{d_0, K_v^{\frac{v+1}{v}}\} \left(\tfrac{1}{\epsilon^2}\right)^{(v+1)} .$$

□

Notice that the bound in Proposition 4.2 holds uniformly over all probabilities on the sphere. This fact allows us to use a version of the entropy integral in the final stage of our argument. Recall Theorem 3.2. Adapted to our current setting, we have the following corollary:

Corollary 4.3. *There exists an absolute constant $K > 0$ such that for any $u > 0$ with probability at least $1 - 2e^{-u}$ for all $W_{x,y} \in \mathcal{W}^{n,s}$:*

$$\sum_{k=1}^{m} \left(d(x,y) - 1_{W_{x,y}}(g_k) \right) \leq K\sqrt{m} \left(\sqrt{u d(x,y)} + \int_0^{\sqrt{d(x,y)}} \sqrt{\log(D(\mathcal{W}^{n,s}, t))} dt \right).$$

We adjust this result in two ways to produce the main results of this section. First, increase the right side of the inequality by replacing all distances with one. Now that the bound is uniform over pairs of signals in $\mathbb{S}_s^{n-1}$, we observe

$$\sup_{x,y \in \mathbb{S}_s^{n-1}} \sum_{k=1}^{m} \left(d(x,y) - 1_{W_{x,y}}(g_k) \right) = \sup_{x,y \in \mathbb{S}_s^{n-1}} \left| \sum_{k=1}^{m} 1_{W_{x,y}}(g_k) - d(x,y) \right|.$$

This is because $1_{W_{-x,y}} = 1 - 1_{W_{x,y}}$ (a.s.), and $d(-x,y) = 1 - d(x,y)$. Thus we have:

Corollary 4.4. *There exists an absolute constant $K > 0$ such that for any $u > 0$ with probability at least $1 - 2e^{-u}$,*

$$\sup_{x,y \in \mathbb{S}_s^{n-1}} \frac{1}{m} \left| \sum_{k=1}^{m} 1_{W_{x,y}}(g_k) - d(x,y) \right| \leq \frac{K}{\sqrt{m}} \left(\sqrt{u} + \int_0^1 \sqrt{\log(D(\mathcal{W}^{n,s}, t))} dt \right).$$

After applying the uniform entropy bounds of Proposition 4.2 in the above corollary and setting $u = \log(2/\epsilon)$, Theorem 2.3 is immediate.

5. The RIP of Noisy 1-Bit Embeddings

5.1. **Noisy RIP with the distorted metric on the sphere.** This section proves Theorem 2.5. We again consider $A \in \mathbb{R}^{m \times n}$ with rows $\{g_k\}_{k=1}^m$ drawn independently from the standard Gaussian distribution $\mathcal{N}(\mathbf{0}, I_n)$. We are now interested in the case of additive white noise prior to quantization; let $\eta \sim \mathcal{N}(\mathbf{0}, \sigma^2 I_m)$. Then the Hamming distance between the images of two signals under the 1-Bit embedding Φ_A^η is

$$d_{\mathcal{H}^m}(\Phi_A^\eta x, \Phi_A^\eta y) = \frac{1}{m} \sum_{k=1}^{m} \frac{1 - \operatorname{sgn}(\langle x, g_k \rangle + \eta_k) \operatorname{sgn}(\langle y, g_k \rangle + \eta_k)}{2}. \tag{5.1}$$

Fix $x, y \in \mathbb{S}^{n-1}$. Let $g \sim \mathcal{N}(\mathbf{0}, I_n)$ and $\mu \sim \mathcal{N}(0, \sigma^2)$ be independent. Then $\binom{\langle x,g \rangle + \mu}{\langle y,g \rangle + \mu}$ is a Gaussian vector with covariance matrix

$$\begin{bmatrix} 1 + \sigma^2 & \langle x, y \rangle + \sigma^2 \\ \langle x, y \rangle + \sigma^2 & 1 + \sigma^2 \end{bmatrix}.$$

A computation similar to 4.1 yields

$$\mathbb{P}\big(\operatorname{sgn}(\langle x, g \rangle + \mu) \operatorname{sgn}(\langle y, g \rangle + \mu) = -1 \big) = \tfrac{1}{\pi} \arccos\left(\tfrac{\langle x,y \rangle + \sigma^2}{1 + \sigma^2} \right).$$

This last quantity is $d^\sigma(x,y)$, defined in (2.1). We'll see soon that d^σ is in fact a metric; this is the distance with which $\mathbb{S}_s^{n-1}$ is naturally endowed in the presence of additive white noise. The object in the (s, δ, σ)-RIP that we aim to bound is

$$\sup_{x,y \in \mathbb{S}_s^{n-1}} \left| d_{\mathcal{H}^m}(\Phi_A^\eta x, \Phi_A^\eta y) - d^\sigma(x,y) \right|.$$

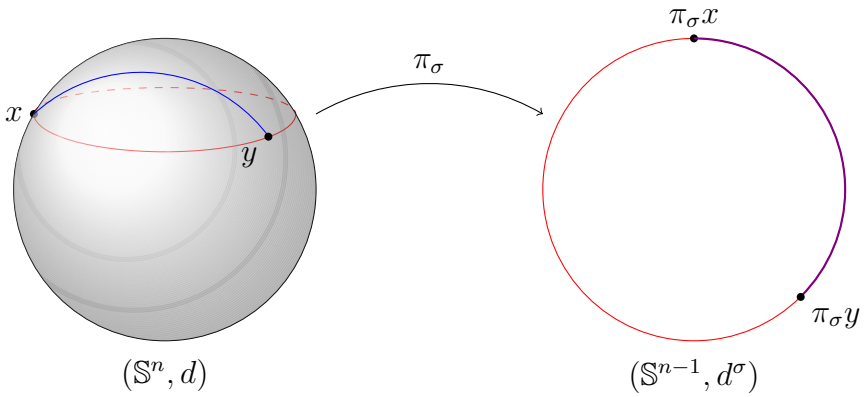

FIGURE 1. If $\pi_\sigma : \{p \in \mathbb{S}^n : \langle p, e_{n+1}\rangle = \sigma\} \to \mathbb{S}^{n-1}$ is the normalization of the projection onto the first n coordinates, then $d^\sigma(\pi_\sigma x, \pi_\sigma y) = d(x, y)$.

Appealing to the methods in Section 4, we rewrite the additive noise as an inner product by increasing the Gaussian measurements' dimension by one and lifting the sphere to one higher dimension by padding with σ^2. Introduce the following notation:

$$x_\sigma = \tfrac{1}{\sqrt{1+\sigma^2}}(x^1, \dots, x^n, \sigma) \in \mathbb{S}^n_{s+1}.$$

Let $h = (g^1, \dots, g^n, \frac{1}{\sigma}\mu)$ and notice $\langle x_\sigma, h\rangle = \frac{1}{\sqrt{1+\sigma^2}}(\langle x, g\rangle + \mu)$ and $h \sim \mathcal{N}(\mathbf{0}, I_{n+1})$. Denote by $W^\sigma_{x,y}$ the wedge in $\mathbb{S}^n$ relative to x_σ and y_σ, i.e.,

$$W^\sigma_{x,y} := H_{x_\sigma} \triangle H_{y_\sigma}.$$

Then $\operatorname{sgn}(\langle x, g_k\rangle + \eta_k) \neq \operatorname{sgn}(\langle y, g_k\rangle + \eta_k)$ if and only if $h_k := (g^1_k, \dots, g^n_k, \frac{1}{\sigma}\eta_k) \in W^\sigma_{x,y}$. The Hamming distance in (5.1) can be reformulated as

$$d_{\mathcal{H}^m}(\Phi^\eta_A x, \Phi^\eta_A y) = \frac{1}{m}\sum_{k=1}^m 1_{W^\sigma_{x,y}}(h_k).$$

Furthermore, notice that

$$\langle x_\sigma, y_\sigma\rangle = \tfrac{\langle x,y\rangle + \sigma^2}{1+\sigma^2},$$

hence $d^\sigma(x, y) = d\left(x_\sigma, y_\sigma\right)$, where we abuse notation to allow $d\left(\cdot, \cdot\right)$ to denote the normalized geodesic distance on $\mathbb{S}^n$; see Figure 1 for an illustration. It is now apparent that d^σ is indeed a metric on $\mathbb{S}^{n-1}$.

This brings our attention to the following object:

$$\sup_{x,y\in\mathbb{S}^{n-1}_s} \left| \frac{1}{m}\sum_{k=1}^m 1_{W^\sigma_{x,y}}(h_k) - d^\sigma\left(x, y\right)\right|,$$

where the row vectors h_k are independently drawn from $\mathcal{N}(\mathbf{0}, I_{n+1})$. At this point, the argument of Section 4 applies so long as we can estimate the VC-dimension of

$$\mathcal{W}^{n,s}_\sigma := \{W_{x,y} \in \mathcal{W}^{n+1,s+1} : x^{n+1} = \sigma = y^{n+1}\}.$$

Since $\mathcal{W}^{n,s}_\sigma \subset \mathcal{W}^{n+1,s+1}$, it is clear that

$$VC(\mathcal{W}^{n,s}_\sigma) \leq VC(W^{n+1,s+1}) \lesssim (s+1)\log\left(\tfrac{n+1}{s+1}\right) \lesssim s\log(n/s).$$

5.2. **Noisy RIP with geodesic metric on the sphere.** The deviation of d^σ from the geodesic distance is exaggerated at antipodes. That is, for any x and y on the sphere, $|d(x,y) - d^\sigma(x,y)| \leq d(x,-x) - d^\sigma(x,-x)$. In what is to come, all suprema are over $x, y \in \mathbb{S}_s^{n-1}$. If one prefers a bound of the form

$$\sup |d_{\mathcal{H}^m}(\Phi_A^\eta x, \Phi_A^\eta y) - d(x,y)| \leq \delta,$$

it is enough for

$$\begin{aligned}\sup |d_{\mathcal{H}^m}(\Phi_A^\eta x, \Phi_A^\eta y) - d(x,y)| &\leq \sup |d_{\mathcal{H}^m}(\Phi_A^\eta x, \Phi_A^\eta y) - d^\sigma(x,y)| + \sup |d^\sigma(x,y) - d(x,y)| \\ &\leq \sup |d_{\mathcal{H}^m}(\Phi_A^\eta x, \Phi_A^\eta y) - d^\sigma(x,y)| + 1 - \tfrac{1}{\pi} \arccos\left(\tfrac{\sigma^2-1}{\sigma^2+1}\right) \\ &\leq \delta.\end{aligned}$$

Corollary 2.6 follows easily from Theorem 2.5 and this observation.

Acknowledgements. The author would like to thank Michael Lacey for suggesting the project and pointing out the utility of the VC theory, as well as for useful conversation on the subject and layout of the article.

References

[1] L. Anitori, A. Maleki, M. Otten, R.G. Baraniuk, and P. Hoogeboom, *Design and analysis of compressed sensing radar detectors*, IEEE Trans. Signal Process. **61** (2013), no. 4, 813–827. ↑1

[2] J. Bobin, J.L. Starck, and R. Ottensamer, *Compressed Sensing in Astronomy*, IEEE Sel. Top. Signal Proc. **2** (2008), 718. ↑1

[3] D. Bilyk and M. T. Lacey, *Random Tessellations, Restricted Isometric Embeddings, and One Bit Sensing*, ArXiv e-prints (2015), available at 1512.06697. ↑2, 3, 6

[4] P. T. Boufounos and R. G. Baraniuk, *1-Bit Compressive Sensing*, in Proceedings of Conference on Information Science and Systems (CISS), Princeton, NJ (March 2008). ↑2

[5] E.J. Candés, *The Restricted Isometry Property and its implications for compressed sensing*, C. R. Math. Acad. Sci. Paris **Ser. I 346** (2008), 589-592. ↑1

[6] Anirban DasGupta, *Probability for Statistics and Machine Learning: Fundamentals and Advanced Topics*, 1st, Springer Publishing Company, Incorporated, 2011. ↑4

[7] Simon Foucart and Holger Rauhut, *A Mathematical Introduction to Compressive Sensing*, Birkhäuser Basel, 2013. ↑

[8] L. Jacques and K. Degraux, *Quantized Iterative Hard Thresholding: Bridging 1-Bit and High-Resolution Quantized Compressed Sensing* (2013). ↑2

[9] L. Jacques, J.N. Laska, P.T. Boufounos, and R.G. Baraniuk, *Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors*, IEEE Trans. Inform. Theory **59** (2013), no. 4, 2082–2102. ↑2

[10] M. Lustig, D. L. Donoho, and J. M. Pauly, *Sparse MRI: The application of compressed sensing for rapid MR imaging*, Magn Reson Med **58** (2007), 1182–1195. ↑1

[11] Dmitriy Panchenko, *Some Extensions of an Inequality of Vapnik and Chervonenkis*, Electron. Comm. in Probab. **7** (2002), 6:55–6:65. ↑4

[12] Yaniv Plan and Roman Vershynin, *Dimension reduction by random hyperplane tessellations*, Electron. Comm. in Probab. **7** (2002), 55–65. ↑2

[13] V. N. Vapnik, *Statistical Learning Theory*, Wiley, New York (1998). ↑

School of Mathematics, Georgia Institute of Technology, Atlanta GA 30332, USA
E-mail address: spencer@math.gatech.edu